plus augmenting device", such that the "combined processor" is able to execute a Minimal P Individual.

Moreover, such targets of education are cost effectively educable. It is surely worth some trouble to secure the necessary preconditions for ethical reasons or aesthetic reasons (mental cripples are remarkably ugly). Recent data on equipment costs indicate that the exercise can be justified, as well, on purely economic grounds.

Chapter 11. Consciousness, Self-Reference, and Some Tentative Relations between the Microtheory and the Macrotheory

The dominant theme, obtrusive amongst a potpourri of topics considered in this chapter, is consciousness and its relation to macrotheoretic indices of uncertainty (H), correct belief (θ) and the like. Since there is a tendency to avoid the subject of consciousness as ineluctable, or even because it is held in disrepute, many familiar examples are cited and, as a result of that, the presentation is longwinded and slightly redundant.

Although much of the argument is relegated to the next volume, states of consciousness have so much pragmatic importance (especially in education) that it is necessary to provide an overview of our position in this matter, if the reader is to be left with even a partly completed story (in which, for example, the microtheory and the macrotheory are rationally connected though not at this stage fully unified).

1. Consciousness and Synchronicity

The occurrence of understanding in a strict conversation between A and B, localised in L processors α and β, implies a synchronisation of procedures executed in α and others undergoing execution in β. It is "obvious" (in a sense) that the execution of these procedures must be synchronised in order that the transactions of understanding shall occur; it is, perhaps, less obvious that in the absence of the systemic form of an understanding, as it is shown in Icon 3 and Icon 5, there is no reason to suppose that procedures executed in α are synchronised with those in β and that if this synchronisation occurs because of the realisation of this systemic form, then there is a transfer of information, i.e. the understanding "is the case" (not just "a simulation of understanding is performed"); or, to put the matter strongly enough to invite criticism, A and B are, under these circumstances, aware.

In fact, if it happens (as it does in the icons) that A and B are separated (in α and β) by an interface and if the sprout of the A, B, conversation is anchored upon R_i then a stronger statement is possible, namely, that

"A is conscious with B of R_i"

which implies the generalisation.

"The A, B, conversation (itself a P Individual) is aware" (under the circumstances where the awareness is observable, it is also possible to assert the topic, R_i, of which at least one participant is conscious, as well).

The matter can be stated concisely, as follows. The systemic form of Icon 3 may or may not be realised, even if all of the systemic prerequisites are satisfied (that is, the specification of a conversational domain and the inscription, as latent or potential operations, of the various Procs). If it happens that there also exist L processors of an appropriate type then the conditions for awareness are realised kinetically i.e. there is awareness, though it may be unobservable; for example, the realisation in one processor in Icon 13. This icon, or any refinement of it, means that there exist in one L processor initially asynchronous loci of control and that the entity A, B, said to be aware, is so because these loci of control are synchronised by information transfer mediated by a conversation internal to the one processor. If, on the other hand, the realisation or actualisation or reification of systemic form takes place at an interface separating processors α and β (Icon 5) then the information transfer responsible for the synchronisation of A and B is understanding, in a strict conversation on domain R, which is observable over occasions, n.

2. External Observers

An external observer of a strict conversation is an individual who is conscious of the conversation as taking place across a given interface with some other observer(s). He communicates with the other observer(s) in the observational metalanguage L*. He is characterised as an external observer insofar as he regards the conversation under scrutiny as an it. In other words, if the external observer is himself represented by an icon, then the interface at which the strict conversation is observed becomes his modelling facility and his interaction with it, during the strict conversation is restricted by the capabilities of an interface processor. In particular, his description is exactly the record made by the recording equipment which appears in each of the

conversational icons. Insofar as the strict conversation takes place over occasions, n, the external observer's stopwatch is synchronised with the execution clock (so that n is an index in the external observer's modelling facility, with exactly the same status as the index τ in the participant's modelling facility. Moreover, the external observer's modelling operations are clocked by an index T which has exactly the same status as the modelling operation index, t, in the participant's modelling facility. Finally, in regarding the conversation as impersonal (an it), the external observer is required to maintain the same (structure/function) distinctions as a participant who regards a modelling facility impersonally.

2.1. Consequently, the external observer is at liberty to make a model (for example, by specifying a CET heuristic, or an uncertainty regulation heuristic) before the conversation when his $n = 0$ (equivalent to $\tau = 0$) and his $T > 0$ (equivalent to $T > 0$). But during the conversation his $T = 0$ and his $n > 0$ and he cannot

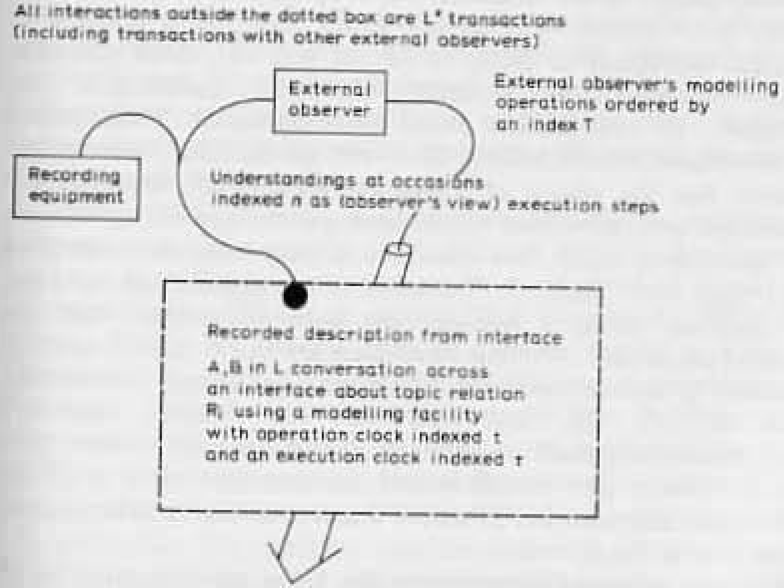

Fig. 11.1. The status of an external observer. If $T > 0$ then $n = 0$. Throughout a strict conversation, the observer refrains from modelling operations as a result of which $T = 0$ and the conversation is impartially observed as the execution over stages, $n > 0$, of a joint model with distinct parts α and β separated by an interface.

change the model without halting the conversation (if his $t > 0$ his $n = 0$. Failing that, he is no longer an external observer). All this is summarised by Fig. 1.

2.2. The external observer's microtheory or molecular theory of an experiment is the icon for that experiment. Thus Fig. 1 is filled out by imposing the appropriate icon upon the blank space labelled "strict L conversation between A and B on R".

3. Macrotheories

The external observer is able to impose several kinds of macrotheory upon the framework afforded by this icon in which the fundamental variates are statistical quantities. It is instructive to consider the possibilities and the interpretations that are legitimately given to the statistical measures.

3.1. *Selective macrotheory.* The least committal macrotheory is founded upon basic measurements that are (frequency) entries in contingency tables obtained by recording the occurrence of events, joint events, and so on. From these, it is possible to calculate entropies, in place of (or as well as) other statistics, which represent the external observer's uncertainty (or, conversely, his information) about the conversation. These indices are uncertainties/informations in a very special sense; namely, the observer has the stance of a selective information theorist or a statistician and subscribes to the appropriate laws of deportment and calculating rules. For example, he may compute quantities like I* and I of Chapter 2. Moreover, he can examine interactions, the simplest being a transmission between entities that are regarded as unitary. Obvious candidates are the M Individuals α, β, separated by an interface (in fact, they may be the only candidates). Thus, if $I(\alpha)$, $I(\beta)$ and $I(\alpha,\beta)$ are the external observer's uncertainties/informations about α, β, and the joint system, then $T(\alpha, \beta) = I(\alpha) + (\beta) - I(\alpha, \beta)$ so that, on these grounds, α and β are statistically dependent, if $T(\alpha, \beta) > 0$ and statistically independent, if and only if $T(\alpha, \beta) = 0$.

Selective information measures like these are computed in L* (restricted to standard statistical usage) and refer directly to the external observer. Additional hypotheses, to the effect that his uncertainty or information images the participant's uncertainty or information, are needed to equate measures of that kind, with

indices of performance or difficulty (for example, in Chapter 2, the notion that ρ is an estimate of I and that η is an estimate of I*).

3.2. *Selective process macrotheories.* One of the statistical precepts which underpins the calculation of selective entropies is that contingency tables span alternative sets (choice sets, conditions that are both exclusive and exhaustive). So, as a result, any selective uncertainty or information value may be interpreted as the amount of *selective work* that an idealised serial dichotomiser or "selective operator" designed as a goal directed system, must do to select one and only one correct alternative (correct to the external observer and to the dichotomiser as well). The word "idealised" means only that no comment is made upon how the selection is performed or, equisignificantly, the selective work is averaged over all possible computing methods.

The value of I is maximised if the alternatives are a priori equiprobable, in which case, for M alternatives, selective work = Log M = I. Otherwise, I has a lower value which may be expressed as the number EM of (imaginary) equiprobable alternatives equivalent to the actual set (biassed by unequal a priori probabilities). That is selective work = Log ME = I. Finally, if no alternatives are specified, I is undefined. The "equivalent alternative" trick is played in calculating the guessing index used in Appendix B, for example.

Hick (1952) noticed a correlation, in choice reaction time experiments, between the reaction time (RT) and the number of choice alternatives. For example, it is possible to write "RT = Constant x Log M" and, in general, "RT = Constant x EM". In this form, Hick's law asserts that if man is viewed as a selective operator, the rate of reduction in selective uncertainty is constant, or (as a converse statement) the rate of gain of information is constant. Thus the amount of selective work is measured by RT, for any correct selection. It should be stressed that selective work carries no commitment about the nature of the selective operation. In particular, this operation may be conscious or not. If not, the operation converges to a correct alternative because it is designed to do so. Under this ruling it is legitimate to interpret the operation time, measured in units of observer's stopwatch time, as an index of the amount of selective work needed to select one out of ME equivalent alternatives, and thus an index of I.

For example, in Chapter 2, these selective work indices are used (in common with selective indices) in computing the value of a performance ρ and thus I. In turn I is interpreted as an indirect estimate of a participant's subjective uncertainty or subjective information. The essential requirements are that a goal shall be specified (some R_i in R) and that participant A shall be restricted to selection from a set that he regards as an alternative set[26].

To obtain this result, in the case when the observed system is conscious, the external observer must truncate operations that would otherwise be explanatory. To do so, he must arrange that whatever agent is introduced as a further participant B, halts the t clock (setting t = 1, so that conclusions or outcomes are selected rather than derived; the expedient of Icon 11 and Icon 12). The measure of selective work does not, in itself, rely upon the existence of consciousness, only its connotation as an indirect estimate of subjective uncertainty or subjective information.

3.3. *Degrees of belief.* If the system is conscious (and a subjective estimate is nonsensical, otherwise) it is also possible to obtain a direct estimate of subjective uncertainty or information from confidence estimates over sets of alternatives obtained (as before) by truncating on explanation (setting t = 1) provided that the interval needed to form the confidence estimate is less than the interval needed to derive the correct outcome or to furnish (though not necessarily to verbalise) an explanation.

3.4. *An indeterminancy principle.* The following digression gives a deeper meaning to the notion of when it is and is not possible to obtain a direct estimate of subjective uncertainty/information. Perhaps the commonest modelling facility is the body (A's body α, for example). It has, of course, the peculiarity

[26] Given a language and knowing how the computation may take place (i.e. by applying inference rules in this language) the alternative set restriction no longer applies and the notion of "selective operator" is replaced by "language processor". Unfortunately, the information theories which can be metricised in this specific way are currently restricted to monadic languages (as in the "semantic" information measures of Bar Hillel (1964), Carnap (1950), and Chiaraviglio (1971) and the inference chains corresponding to explanations are not, at this stage of development, of full psychological value. Notice, however that Bruner, et al. (1956) use "information" in (roughly) this sense; "number of possible inferences excluded by a statement", so does Harrah (1963). Seigman and Stapleton (1971) have also applied such measures with success to psychological processes.

that neither an external observer nor his agent, B, is able to inhibit the execution of a model. This peculiarity is often (and rightly) called autonomous action. For example, the novice, learning to typewrite and told to "press the key of s" makes a model in his neuromuscular modelling facility which contains his current knowledge of how to press "the key of s". Unless very odd precautions are taken, this model, right or wrong, is executed immediately that the first operation is defined. The execution is autonomous. If the participant is, at the same moment, asked to give a confidence estimate expressing his degree of belief, that one or other key is the "key of s", his attempt to do so is abortive if the autonomous action selects the "key of s" before he can state his degree of belief.

The phenomenon pinpoints a basic indeterminancy of psychological observation. If there is time, the direct estimate is the best estimate of subjective uncertainty or information; if there is not, the best estimate is the correct response latency or *selective work*. When there is not time to obtain a direct estimate, specificity is lost.

3.5. *Direct estimates.* For example, H and θ (Chapter 2, Chapter 4 and Chapter 6) are direct estimates evaluated in this manner, over the alternative answer set of a $\underline{\text{PQuest}}^0$. Clearly, the alternative set need not span the immediately considered topic relation (the look ahead uncertainties H^* and θ^* do not do so), nor need it refer to alternative outcomes such as solutions to an L^0 problem (for instance, $\underline{\text{PQuest}}^1$ alternative sets, used in the calculation of strategic uncertainties, span alternative classes of nodes). However, the present discussion is confined to confidence estimates on the alternative answer sets of $\underline{\text{PQuest}}^0$s and, for the most part, to $\underline{\text{PQuest}}^0$s spanning a subgoal topic relation in the sprout of a strict conversation, only.

4. Subjective Uncertainty

The crucial point is that macrotheories are all based on measurements over sets of alternatives or choice sets; but the idea of a "choice" is artificial, to the extent that someone in the position of an external observer truncates an explanatory or modelling operation that might otherwise have taken place, thus constructing the alternative set (canonically, the external observer does so by specifying that t = 1).

To avoid confusion, let us use the words doubt/belief to designate the subjective uncertainty experienced by a participant and the words uncertainty/certainty to designate direct estimates of doubt/belief obtained by eliciting a confidence estimate over an alternative set and calculating an information measure.

4.1. *Belief and awareness.* Doubt/Belief is a property of awareness. If the quantity is undefined for participant A then A is unaware. If it is defined in the context of an alternative set, spanning a topic relation R_i, then A is aware of R_i and may be conscious of R_i with B. The discussion is restricted (the restriction is illusory, in fact) to a modelling facility where models are built which, on execution, bring about R_i. Further, some of the discussion is restricted to the case of a one clocked modelling facility (this restriction is genuine, it is relaxed to produce further insights).

Suppose you are the participant. Your doubt/belief has at least three dimensions (Fig. 2). The first of these, d_0, is a doubt about what you are attending to. If you (A) are attending to R_i (or, in particular, if you are conscious of R_i, in a strict conversation with B) then this component of doubt is, or is nearly, zero; $d_0 = 0$. If and only if that is so, can you specify a set of alternatives that span R_i (that are exclusively disjunctive L statements which you believe may satisfy R_i). Nothing is lost by supposing you are at

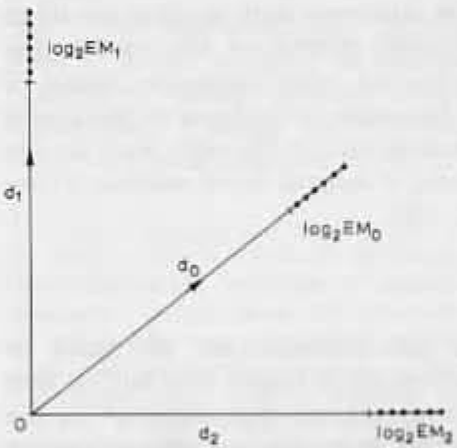

Fig. 11.2. The coordinates of doubt as it is experienced by a participant. Dotted regions on the coordinates are undefined; doubt is greater than a value on the ME (equivalent alternatives) scale since there is no alternative set.

liberty to specify these statements of your own accord, rather than being given the alternatives (as you would be if you were posed a standard PQuest[0]). In general, if you are asked by B or me to specify alternatives in respect of a modelling facility in the context of R_i, you will specify methods of modelling R_i in such a way that the execution of the model would bring about R_i. The modelling facility might be your own brain and/or it might be a one clocked modelling facility like STATLAB.

4.2. *Types of alternative sets.* If, on the assumption that $d_0 = 0$ and that you have specified a set of alternative methods, you are asked (by B or me) to assert your belief or doubt regarding these methods at some occasion n (and, recall, the external observer can only order your activities as a sequence of occasions, n), then there is an ambiguity that may be resolved by qualifying the question in a manner that makes sense to you but has no metrical significance to me (because I look at your activity as ordered by occasions n, the ambiguity is irreducible to me). In other words, the question that elicts an expression of doubt/belief may be rephrased as, two questions, that are, from your point of view, distinct; though they are both open to several equisignificant verbalisations; for example, as follows (all of them to be answered either deterministically by the choice of one alternative or by a subjective probability distribution over all of the alternatives in the set).

(1) "How did you model relation R_i"? or "What method (of modelling R_i) is used"? or "Which description of R_i do you prefer"?

(2) "How will you proceed in modelling relation R_i"? or "What action comes next in modelling R_i"? or "Which option (for modelling R_i) do you prefer"?

Because of the temporal polarity manifest from your point of view within the occasion, call your doubt or belief in respect of the (M_1) alternative methods that you specify as possible methods in giving an answer to (1) a retrospective doubt, d_1. For the same reason, call your doubt in respect of the M_2 alternative options you specify as possible options in answering (2) a prospective doubt d_2. The external observer (me), or you for that matter, can estimate the value of d_1 and d_2 if and only if $d_0 = 0$, when you can list alternatives or interpret a list you are given.

One peculiarity of d_1 and d_2 is that their highest values are undefined and discontinuous in that region. To see this, notice that the value of d_1 depends upon M_1 and your degrees of belief about these M_1 alternative methods and the value of d_2 depends upon M_2 and your degrees of belief about these M_2 alternative options. Using the trick of equivalent equiprobable alternatives, d_1 may be expressed as a function of the number of equivalent equiprobable alternatives EM_1 (which is zero only if $M_1 = 0$) and d_2 as a function of EM_2 (which is zero only if $M_2 = 0$). But the graphs of d_1/EM_1 and of d_2/EM_2 are discontinuous since d_1 may be high valued either because EM_1 is large or because $EM_1 = 0$ and there are no methods; (similarly) d_2 may be high valued either because EM_2 is large or because $EM_2 = 0$ and there are no options. The essential constraint upon the system is that if $EM_1 = 0$ then $EM_2 = 0$ and vice versa; (if d is high because there are no options and vice versa, this is a direct consequence of the fact that options are prospective methods; conversely, if there are methods then there are options and if there are options then there are methods).

If the question is not qualified (as a means for distinguishing d_1 and d_2) you are still aware of belief. Similarly, if $d_0 > 0$ you are either unaware or still aware of an unformulated amorphous belief. Call this sense of becoming *prescience*; it is conversely expressed by a mixture of d_0, d_1 and d_2.

4.3. *Specific instances of doubt and belief.* Before going further, these ideas bear exemplification. I have sketched, in Fig. 3, a few of the many commonly experienced doubt/belief patterns, on the assumption that $d_0 = 0$ so that ME_1 and ME_2 are specified to give actual scale length to the co-ordinates of d_1 and d_2.

(a) d_1 high, d_2 high: Prescience of R_i, or ignorance of exactly what methods bring about R_i, or what options to take in modelling.

(b) d_1 low, d_2 low: Absolute lack of doubt about R_i.

(c) d_1 undefined, d_2 undefined: Ignorance of anything. You are not aware.

(d) d_1 low, d_2 high: Certainty of a method for modelling R_i but doubt of whether the model will bring about R_i, on execution.

(e) d_1 high, d_2 low: Doubt about how R_i is modelled, but certainty over the choice of option.

(f) d_1 high, d_2 high: Some doubt about the method of modelling R_i and also about the next option.

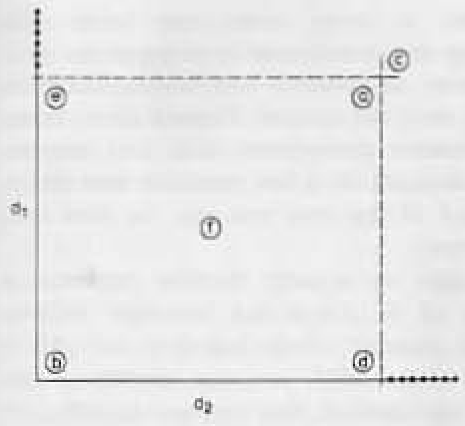

Fig. 11.3. A participant's d_1, d_2 coordinates drawn on the assumption that $d_0 = 0$. Points indicate familiar and important experiences; with the exception of point (c) the participant is aware of R_i (the topic relation spanned by alternative sets) and may be conscious of it. If so, his beliefs/doubts are properties or qualities of consciousness/awareness of R_i.

The interpretation of (f) is equivocal. If you are engaged in modelling a relation R_i then your options constitute a set of alternative modelling methods and your doubt is engendered by choosing one advancement of the model which is to be executed. It is also possible to present (degenerate) options that are the set of results or outcomes achievable by executing the model. For example, the questions $\underline{\text{PQuest}^0}$, so far considered, are degenerate, though they need not be. The alternative answers to $\underline{\text{PQuest}^0}$ constructed by the rules in Appendix I are options reduced to outcomes. This state of affairs is secured by setting $t = 1$, and $\tau = 1$, and thus truncating a modelling operation or explanation.

4.4. *The participants brain and body as a modelling facility.* Consider the case in which your brain and body figures as the modelling facility (not generally one clocked).

Condition (e) in Fig. 3 corresponds to a familiar and often occurring experience, in respect of an overlearned skill like typewriting (for the professional typist) or motor car driving (for the professional motor car driver). If you occupy one of these roles and the topic relation R_i bears upon your occupation then you are certain of what to do and thoroughly uncertain of how you do it; for the simple reason that there are many concurrently executed procedures (models in your brain and body qua modelling facility). In treating the procedures as alternatives you are required to conceive them as parallel but not concurrent procedures and to prefer one over the others. Viewed thus, there are so many (as though parallel) procedures that you cannot disentangle them, or give an account of a few plausible and likely methods (still less, an account of the *one* you use, so that the requirement does not make sense).

Condition d of Fig. 3 images an equally familiar experience (Scott and I found reports of it ubiquitous amongst trainee teleprinter operators and also amongst clerks learning new office routines). On first acquiring a new skill you are certain of the method (so far there is only one method, but you are doubtful of its result). The proficient operative has many methods but cannot describe any of them.

As a final example, the common experience of recall from long term storage in the brain (a reconstructive memory operation) may be characterised by condition (d) (you know what recall method you used, but do not know the name, say, of the place or person you are recalling); or by condition (e) (you know the name as a flash of insight but are ignorant of how you reconstructed it because there are so many reconstruction methods you used concurrently); or by condition (f) of Fig. 3 (characterised by both kinds of doubt).

4.5. *Restrictions imposed upon doubt and belief of a simple modelling facility.* Consider the special case of a one clocked modelling facility such as STATLAB (or equally, the case in which you have to give one description or explanation, as in writing it on paper). Suppose that you are asked to express your prospective doubt and your retrospective doubt whilst making a model in a one clocked modelling facility (in response to $\underline{\text{Comm}}^0$) or to furnish one (the one) explanation (in reply to $\underline{\text{EQuest}}^0$). Under these circumstances, there is no retrospective doubt (your model or your explanation up to this point may be entirely mistaken; but you are sure of it). Hence, the act of learning to model or explain R_i is represented by a trajectory of the kind shown in Fig. 4. Starting from ignorance (no method, high d_1; no idea of what to do; high d_2), you build some model. But as soon as you have any model (since you can only have one) d_1 becomes low valued though d_2 may still be high (for your one model may not work to

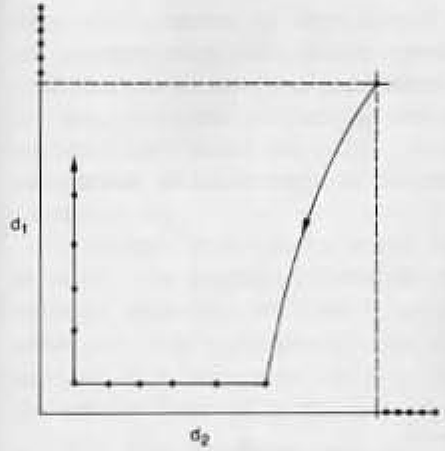

Fig. 11.4. The trajectory or cyclogram traced out by learning (a topic relation R_i) as sampled by experiences of d_1 and d_2. This sketch is characteristic of learning in which one and only one explanation of R_i is required. It is assumed that $d_0 = 0$ so that alternative sets spanning R_i may be specified.

bring about R_i and you may be unconvinced that it will do so). As the model is perfected, d_1 decreases and your conviction increases; and learning to explain R_i in a one clocked modelling facility forces the trajectory through a point (Low d_1, Low d_2). After that, the value of d_1 may increase as more methods of modelling are learned and these may be listed as alternative models. In a many clocked modelling facility (including your own brain and body) other learning trajectories are possible (Fig. 5).

4.6. *The doubt and belief of holist and serialist participants.* Both Fig. 4 and Fig. 5 are drawn to represent an L^0 modelling process or an L^0 explanation, sampled periodically by questions of the type $\underline{\text{PQuest}}^0$, to elicit d_1 and d_2. It is also possible (though the matter is not pursued) to sample an ongoing L^1 modelling operation or an L^1 explanation by $\underline{\text{PQuest}}^1$ (namely to sample doubt or belief about a learning strategy). The figures are the same. But given this interpretation it is worth commenting that the trajectory of Fig. 4 is a serialist trajectory and the trajectories of Fig. 5 are holist trajectories. Due to the one clocked form of STATLAB only the trajectory Fig. 4 can be realised in the current version of CASTE at the level of L^0 modelling under one subgoal (this limit is being eliminated). However, even in the current

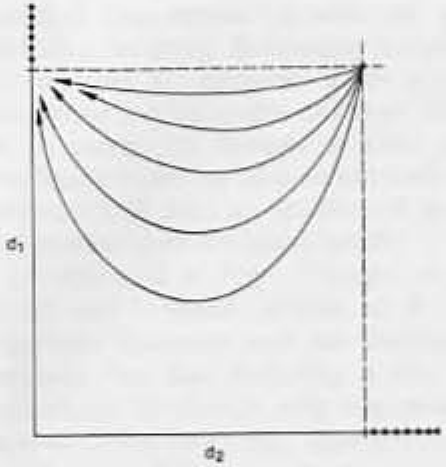

Fig. 11.5. Trajectories or cyclograms characteristic of learning in which the requirement of one and only one explanation is relaxed and methods are allowed to proliferate as the participants doubt about options is decreased by learning.

version of CASTE, Fig. 5 can also be realised in terms of L^1 interactions across the interface.

4.7. *Direct estimation of doubt and belief.* If $d_0 = 0$ (so that a participant A is concentrating upon a particular topic relation, R_i) alternatives spanning R_i may either be elicited, or given by the external observer. A confidence estimate over these alternatives, which is usually obtained as a BOSS measurement (Chapter 4 and Chapter 6), is a direct estimate of prospective belief in each of the alternatives and the uncertainty/information function H (Chapter 2, Chapter 4 and Chapter 6) is a direct estimate of d_2. Moreover, since R_i is known and since the alternative set spans it, a correct belief measure such as θ is a direct estimate of correct belief (the converse of d_2).

These interpretations hold true for confidence estimates elicited by $\underline{\text{PQuest}^0}$ constructed in the manner of Appendix I and also for L^0 which-questions (that are "open-ended" in the sense of allowing for student stipulated alternatives) provided either (a) one and only one of the stipulated alternatives can be recognised as correct, or (b) if more than one stipulated alternative is recognised as correct, then the stipulated alternatives are divided into subsets (perhaps with repetitions) so chosen that one and

only one member of each subset is correct; confidence estimates are elicited over each subset (labelled 1, 2. . . .); an uncertainty/information function is calculated for each confidence estimate as $H_1, H_2, \ldots$; and the several values of the functions are averaged to form $H = \text{mean}(H_1, H_2, \ldots)$. The correct belief function is calculated, as an average, in the same way; so that its value is $\theta = \text{mean}(\theta_1, \theta_2, \ldots)$.

It is also possible to elicit direct estimates of retrospective doubt d_1 (by posing a differently qualified $\underline{\text{PQuest}^0}$) and to obtain distinct estimates of both quantities. Such measures have been obtained, and underpin the rather dogmatic argument in the last section, but they were not used during the experiments (learning probability theory) in the main example.

4.7.1. The sampling scheme, employed for the uncertainty regulation heuristic in CASTE, also contains a "look ahead" uncertainty H* and a "look ahead" correct belief θ*. Both of these pairs of quantities (H, θ and H*, θ*) are calculated from confidence estimates over the alternative answers to a $\underline{\text{PQuest}^0}$ and in both cases there is an assumption that $d = 0$. Finally, in both cases, the sampling scheme can be elaborated to directly estimate the d_1 component of doubt as well as the d_2 component of doubt. So far as the look ahead uncertainties are concerned, the distinction has a proven practical utility.

4.7.2. The ossature of the scheme is shown in Fig. 6. H* and θ* are calculated at the aim node from a confidence estimate over the answers to $\underline{\text{PQuest}^0}$ (Aim). H and θ are calculated at the subgoal nodes from a confidence estimate over the answers to $\underline{\text{PQuest}^0}$ (Subgoal); (They are also averaged over all subgoals in the workset of a goal).

Suppose that $\underline{\text{PQuest}^0}$ (Subgoal) is posed. Under a CET heuristic, or an uncertainty regulation heuristic, if the student changes his attention he may legally attend to any topic relation with a node in the Ent Set of aim (provided this node is not marked as being understood). Thus, if $d_0 > 0$, there is a set of alternatives over which d_0 may be estimated; namely this legal set. The uncertainty regulation heuristic is based on the idea that the Aim node is the most (subjectively) probable member of a look ahead set (of nodes of appreciated topic relations i.e. describable but not addressable topic relations, in the sense of Chapter 4).

In this respect, the one aim restriction upon the existing system

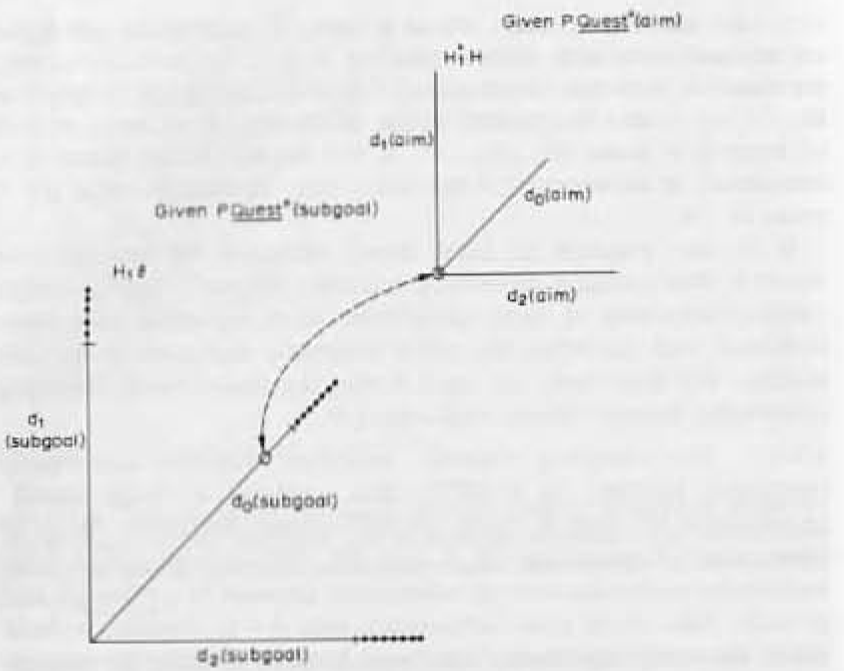

Fig. 11.6. Outline of the uncertainty measures used in a conversational system indicating one subgoal node for $d_0 = 0$ and one aim node.

is an embarrassment; there may be more than one most (subjectively) probable or most favoured topic relation (i.e. the appreciated set is not necessarily one alternative set but several of them). Notice that a serialist student, for whom goal = subgoal = aim, has no measurable uncertainty about what to attend to until the currently addressed topic relation (with node = goal = subgoal = aim) is marked as being understood. A holist student, on the other hand, has a measurable uncertainty about what to attend to, if he is allowed to exhibit it. In either case, the estimate of this quantity is elicited by posing an L^1 question (the simplest form is a $PQuest^1$) and the clutch of method/option uncertainties obtained from confidence estimates over the answers to L^1 questions are the strategic uncertainties mentioned in Chapter 4 and Chapter 6.

With the restriction to one aim node (which is deprecated) it is possible to represent the aim node as lying at a certain distance along d_0 (subgoal) even though the student is operating under an

L^0 command that focusses his attention (apart from any uncertainty sampling L^1 questions) upon the subgoal topic relation. As the serialist student has no uncertainty (because he has no alternative-set of next loci of attention (aim = subgoal) it is immaterial whether his aim node is located beyond the maximum value of d_0 or at the zero value (aim is subgoal). For the holist student, the locus of the aim node along d_0 is (as shown in Fig. 6) a point between zero and the maximum value of d_0.

4.7.3. Suppose that $PQuest^0$ (Aim) is posed to elicit confidence estimates from which to calculate look ahead uncertainties. The exercise is trivial in the case of a serialist student if he literally chooses aim = subgoal; the exercise is not trivial in the case of a holist student or a serialist student, who deviates from this strict (but empirically very common) pattern.

The gross direct estimate of H^* decreases in value (and θ^* increases) as the student learns. For the deviant serialist, the pattern resembles Fig. 4. For the holist student the pattern resembles Fig. 5; methods are accumulated and held in mind whilst the student's doubt about options is decreased. The decrease in doubt about options is not due, in either case, to studying the Aim topic relation directly. It is due to learning the one or many subgoal topic relations in the workset of goal. The deviant serialist uses this knowledge to muster essentially one method whereby he may later tackle the explanation of the aim topic (at the point when the current aim becomes a subgoal). Usually, even the deviant serialist student does not deviate a great deal (he has a one goal workset) so that this observation is in no way surprising. The holist student, who generally has several members in the workset of his goal, accesses topic relations of these nodes as subgoals in a series, but he combines the information from all of them in decreasing his uncertainty about selecting amongst the options spanning the topic relation of aim and the existence of such a plurality of methods (Fig. 5) is due to combining the partial methods learned under different subgoals.

Although it is possible to go through the ritual of forming confidence estimates and calculating information functions as direct estimates of doubt and belief, the operation is onerous (from the student's point of view) even using an equipment such as BOSS. It is also quite easy to obtain estimates of unspecific uncertainty which hardly burden the student at all (he is asked periodically, to consider and state his uncertainty about solving problems of the kind encountered under the current topic relation (R_i) and to consider his uncertainty about the methods to employ in doing so). Provided these statements make sense to the student (as they do) and provided the necessary preconditions are satisfied ($d_0 = 0$; set of alternatives over the topic relation R_i), so that the ritual calculation could be carried out, there is no obvious objection to using the unspecific estimates as they stand.

5. Forms of Awareness distinguished in terms of the macrovariables

This section pulls together previous ideas; it completes a number of outstanding arguments by reference to the macro theory and it contains the first steps needed to establish a useful relation between the micro theory (as depicted by the icons) and the macro theory. Although the doubt/belief variables (d_0, d_1, d_2) are used for expository purposes, it is always true that their values may be estimated, either directly (Section 1.3.3) or by a selective work index (Section 1.3.2), and the properties under discussion can thus be evaluated (with more or less assurance) and later quantified.

5.1. *Learning.* Suppose that the execution of any procedure influences the participants doubt, and, generally, that a synchronisation between loci of control, engendered by internal-to-the-system information transfer, reduces this doubt. If so, the micro (or molecular) theory predictions can be mapped onto the macro (or molar) theory.

In any icon for which the P Individuals (participants A, B) are executed in appropriate L processors, the strict L conversation is anchored on R and its sprout rests upon some R_i in R. The "need to learn" condition (Chapter 5) implies an activity in both the L^1 and the L^0 (horizontal) explanation loops that is causally coupled in A and causally coupled in B (the vertical loops). The condition implies that A does things and learns. The latter activity means either the construction of one L^0 procedure ($\underline{\mathrm{Proc}}^0_A$ ij) which brings the representative point (d_1, d_2) to locus (b) in Fig. 3, or the accretion of further L^0 procedures that are members of the L^0 repertoire ($\underline{\mathrm{Proc}}^0_A$ i) also capable of bringing about R_i. Unless R_i changes, there is no possibility of (d_1, d_2) moving rightwards from (b) in Fig. 3 (unlearning R_i). But if one procedure is learned then

others will be, and the retrospective doubt (d_1) will thus increase, moving (d_1, d_2) upwards in Fig. 3. But suppose that d_1 does not increase; either because no more procedures for bringing about R_i can be constructed, or because the participant is told that only one procedure is acceptable (as in a rigid form of instruction). If so, the representative point remains at (b) of Fig. 3.

Using a CET heuristic, of course, an understanding is detected, a fresh occasion begins, and topic relation R_i is changed to some other topic relation R_j.

Failing that (and R_i persisting) procedures are executed but, contrary to the initial assumption, do not change the locus of the point (d_1, d_2). This absurdity is avoided by noting that under these circumstances the participant learns some other relations, willy nilly (perhaps a relation the external observer deems irrelevant). The first step in this transformation is an increase in d_0 (doubt about what to attend to). The next step, selection of some other R_j, reduces d_0 but with respect of R_j (not R_i). In the case of a CET heuristic and any heuristics based upon it, the first step is reduced to a transient. Another R_j is selected without much doubt about what should be attended to. This is the macro theoretic rationale of the uncertainty regulation heuristic in Chapter 6; it is also, from an external observer's point of view, a means of ensuring relevance and measurability.

The situation is summarised in Fig. 7 where I have sketched an underload region (as in Chapter 2) which is prohibited. A

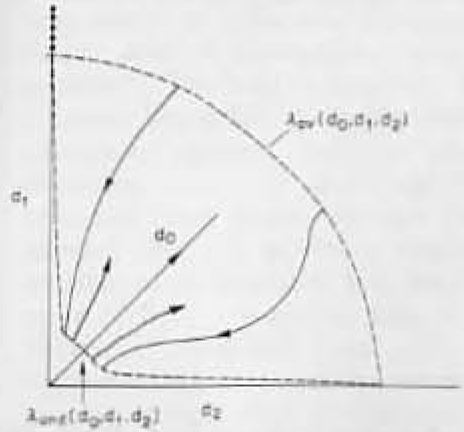

Fig. 11.7. Underload and overload regions constructed as areas of "too low" and "too high" a doubt together with typical learning trajectories.

participant must have belief/doubt about what to attend to (d_0) about method (d_1) or about option (d_2); any or all of them. But (d_0, d_1, d_2) may not remain in the underload region Λ_{und} (d_0, d_1, d_2) of Fig. 7, if A maintains awareness (that is, whilst A, qua P Individual, is executed).

5.2. *Selective work and temporal polarity.* Regard the co-ordinates d_1 and d_2 as monotone functions of EM_1 and EM_2, which may be transformed into selective work co-ordinates (say E_1, E_2) where processing time is measured in terms of the (necessarily many) L processor clock times (Fig. 8). It follows that there is a finite minimum specious present from which E_1 (corresponding to retrospective doubt, d_1) extends, from the participant's point of view, into his past and E_2 (corresponding to prospective doubt d_2) extends from the participant's point of view, into his future. The expression "selective work" is an average over (or a glossing of) the necessarily many clocks in an L Processor, which represents the many branched executions of actual procedures to obtain a describable, but all the same spurious, homogeneous past and future.

5.3. *Underload region.* The measure, E_0, (corresponding to d_0) is only defined if there are strategic alternatives, and E_0 represents effort spent in selecting amongst them to change the locus of the participant's attention. If $E_0 = 0$, then the underload region expressed in the dimensions E_1, E_2 (which is sketched in Fig. 8 as $\Lambda^*_{und}(E_1, E_2)$) corresponds to the region $\Lambda_{und}(d_0, d_1, d_2)$ (in Fig. 7). If $\Lambda^*_{und}(E_1, E_2)$ does not vanish, it follows, from the last paragraph, that E_1, E_2, have temporal interpretation to a participant or P Individual; that any participant who is aware (or, loosely, who is conscious) has a finite specious present and that his minimum value exceeds $\Lambda^*_{und}(E_1, E_2)$.

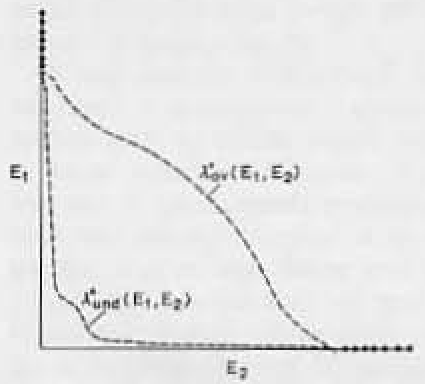

Fig. 11.8. Selective work indices E_1, E_2 (corresponding to d_1, d_2 with $d_0 = 0$) and the underload region (minimal specious present) together with the overload region (the maximum span).

5.4. *Overload region.* By the same token, there are overload regions expressed in terms of doubt (sketched in Fig. 7 as Λ_{ov} (d_0, d_1, d_2)) or expressed in terms of selective work (sketched on Fig. 8 as $\Lambda^*_{ov}(E_1, E_2)$).

$\Lambda_{ov}(d_0, d_1, d_2)$ is a limit imposed upon the mean "size" of a participant represented as his span of awareness. $\Lambda^*_{ov}(E_1, E_2)$ is the mean "size" of the participant represented as the computing activity needed to reproduce the P Individual he is in a fixed context (attending to R_i). However, the computing activity is averaged, over many concurrent loci of control, or L Processor clocks.

5.5. *Limits of awareness and doubt.* Why should there be an overload limit? The real condundrum of psychology is not really why people are conscious or aware, but why there are unconscious processes and events (like the psychological housekeeping of autonomous motor activities) of which people are unaware.

Empirically, that is the case. My conjecture in the matter is as follows. d_1 and d_2 become high valued if EM_1 and EM_2 are large (and they might be indefinitely large). But d_1 and d_2 also become high if $EM_1 = 0$ and $EM_2 = 0$. Under what circumstances could that occur (i.e. under what circumstances could it be impossible to define sets of alternatives, over the procedures that are in progress?). The limit is linguistic. Beyond a certain complexity, it becomes impossible to approximate and speak of an inherently concurrent process, such as class of the procedures under execution in π^1_A, π^0_A, as though it consisted in alternative serial processes (that might, perhaps, be executed, simultaneously, in parallel). But this gambit is required by the rules set out in the calculus of uncertainties and, when the gambit is unplayable, the value of $EM_1 = 0$ and of $EM_2 = 0$. It is this (rather than some arbitrary restriction on L processor capacity) which fundamentally, and perhaps usually, specifies the overload limit. Because of that, there is a limit to awareness. Because of that, in certain special conversational domains, participants are at an advantage in attending to one thing at once.

The limit (or, more accurately, "these limits" for there are several of them) are considered in the next volume and the following clutch of remarks are merely intended to consolidate the known-to-be-incomplete argument by referring it to some everyday experiences and some laboratory situations in which attempts are made to estimate the values d_0, d_1 and d_2.

5.6. *Enforced underload.* Crossman (1956), one of the first psychological information theorists, studied the mentation of cigar rollers. Their work is monotonous but highly skilled, and they are paid to devote some attention to it. Thereby, the working situation imposes, so far as possible, the condition $d_0 = 0$. The work is self paced so that d_2 can be varied by changing the rate of operation; each cigar "poses a problem" but the variation is restricted by the sheer mechanics of rolling a cigar on the thighs. However, many cigars are rolled and, however variable the tobacco leaf, the problems posed by the task do not engender much doubt d_2. Thus, as predicted, d_1 increases; a fact that is manifest in the invention of new rolling methods, some of the utmost detail involving great niceties of movement.

The reader is invited to join me in exercise running "on the spot" for which there are several published tables showing the k/cal/min equivalents of different sustained running rates. The infuriating, distasteful, but salutory skill of "running on the spot" (it does at least have the beneficial effect of waking you up) requires more acumen than might be imagined. In the absence of special equipment, you need to run and maintain balance at a roughly constant pace for a period of 20 or 25 mins. So you count paces and consult a watch placed in front of you. It is easy to loose count and wise to introduce the rule that if you loose count then you take the last recalled figure and start from there (a minor disaster since it prolongs an unpleasant activity). By the nature of the task, d_2 is low. The "problems" of running are not hard to solve. Further, in order to keep count and monitor your limbs, d_0 must approximate zero. Initially, as you start training, d_1 becomes quite high, for you invent enumeration and packing methods, that are different ways of counting paces. But rather few methods fit the ritual and after a few weeks this source of variety is exhausted, so d_1 is low. At this stage, or slightly before it, you experience an almost irresistable need to think of other things; of course, you do so and of course you lose count in the process

which is duly penalised by the rule; so there is a strong incentive to concentrate on running. Performing this task furnishes one of the most clearcut examples of fluctuating attention (attending to some other thing since $d_1 = 0$ and $d_2 = 0$ so d_0 must increase in order to maintain awareness). On the other hand, effective performance depends upon $d_0 = 0$ (unlike knitting or copy-typing the attention cannot be effectively divided). In practice, d_0 fluctuates wildly, which is predictable.

For a clerical or checkout task it is often required that the operation is performed rigidly according to just one method. The chief good reason for this otherwise outrageous demand is a real need to check the checkout and it is much easier to supervise an operation governed, for example, by a flow chart. The embargo upon invented methods (the discouragement of inventors and deviants) ensures that d_1 is low. If the operation is crucial, various expedients are adopted to keep d_0 low. Depending upon the problem complexity, d_2 fluctuates, but not under the clerk's control. If d_2 is accidentally high, then all is well. If d_2 is low, either there is a switch of attention (preceded by an increase in d_0) or the clerk becomes unaware. Sometimes he literally goes to sleep; more commonly, he enters a state of vacuous, not imageful, reverie).

As a final example, consider the word repetition experiments due to Evans (1967) and to Lilly (1961). The respondent sits in a soundproof room and listens to a repeated word; Say the word "Captain"; it is always the same word. Its utterance is recorded on a magnetic tape loop. d_1 is low (the respondent is merely instructed to listen and write down what he hears) and d_2 is low (the input word is, in fact, the same word). The phenomenon observed is an increase in d_0 and a shift to another word (say "Arcane" or "Chaplain"). As a matter of fact, these transformations are fairly regular and may be predicted, for example from Deese's (1965) arguments. But the respondent writes out transformed words as the words he really heard. Hence, there is a fluctuation of d_0 due (by hypothesis), to a low d_1 and a low d_2. The fluctuation has the form

Word believed to be heard	
Captain	$d_0 = 0$
	$d_0 > 0$
Captains	$d_0 = 0$
	$d_0 > 0$
Chaplains	$d_0 = 0$
	$d_0 > 0$
Arcane	$d_0 = 0$
	$d_0 > 0$
Napkin	$d_0 = 0$

and so on.

The experiment has a mildly soporific effect on respondents even in the form described: if d_0 is held low, by "correcting" the transformations (Pask, Scott, Watts et al., Tech. Rep. 1973) the respondent literally falls into a stupor.

5.7. *Enforced overload.* It is quite difficult to achieve the overload condition. Repetitive inputs pose trivial problems, so that high rates of presentation fail to do the trick. Even a random but complex source of inputs may be ineffectual, because the respondent, even knowing the source is random, invents spurious rules that govern it (i.e. he entertains "superstitions" that it is non-random). Somewhat similar comments apply to attempts to fractionate the respondents psyche by irrationally changing his goals (d_0 is high). About the only way to secure overload (Lewis and Pask 1961, Pask 1964) is to control the environment so that any "superstitious" hypothesis is mechanically recognised as a regularity the respondent has appreciated and evidence for such a regularity is eliminated from the subsequent input. In other words, the controller recognises what the respondent counts as a regularity and removes it, selectively. For a rather low rate of removal of regularity, d_0 increases and the controller was used to encourage innovation. But it is possible to remove putatively comprehensible regularities at a rate that nips any hypothesis in the bud; if this condition is employed, the respondent becomes unconscious and falls into a state of reverie.

5.8. *Symbolic and physical operations on limits of doubt and belief.* It is opportune, at this point, to emphasise that the participant (student or whoever) is a P Individual executed in some M Individual, without committment to which M Individual this L Processor is. As a matter of convenience and convention, we are all caught up in a personal-processor oriented view of ourselves; that "I" live in "my" brain ("my" M Individual). There is nothing wrong with this orientation, quite the contrary, provided it is recognised to be a specialised and sometimes fallible point of view. But it is so deeply ingrained in the culture that an appreciable amount of re-thinking is needed to dispel the impression that the orientation is sacrosanct (or even worse the "one truth").

Thus, it is easy to accept the idea that "my brain" might be used as an L processor by "your brain" (meaning alternatively, that I, as a P Individual, interact with you as a P Individual in an L Processor composed of two brains). But equally common-place ideas such as "I am used by an organisation" (which has no specific or localised processor) or "by an idealogy" or that "I am controlled by the tacit norms of a society" (Vickers 1970) or "form part of an appreciative system" (Vickers 1972) have an air of strangeness about them, due entirely to the personal processor orientation. By the same token, the comment (in Chapter 5 Section 2.5) that many of the storage positions needed to execute me as a P Individual are situated in "my environment" rather than "my brain" rings true as a piece of fanciful or figurative imagery, but is hard to accept as a fact (and the comment is intended as a statement of fact). The point is, the whole distinction between "recall" and "recognition" (in the context of which this comment was voiced) depends upon a presupposed personal processor orientation.

In much the same way that we presume "ourselves" spatially located in one L processor we also presume the temporal continuity of consciousness; and for much the same reason. Many activities depend upon just this supposition (driving along a motor-way or flying an aeroplane, for example). It is only when there is an occasion to doubt the continuity of consciousness (or to be technically precise, of awareness) that we realise what a fragile thing this "stream of consciousness" really is.

These strictures are apposite in connection with conditions of doubt/belief that fail to satisfy the requirements on an operating region (when consciousness is lost) and they provide the background in which the next few sections, about self consciousness and the like, make sound sense (or factual nonsense, as you please) rather than vapourings designed to beguile the intellect.

The conditions of doubt/belief that are not compatible with consciousness are chiefly discussed in the next volume. As a preview, we note that a participant can become unconscious (or, as later, unselfconscious) (a) because of operations performed on "his" processor ("his" brain) or (b) because of operations performed upon "his" actual L processor (involving large chunks of "his" environment) or (c) because of symbolic operations performed on him as a P Individual.

Of these categories (a) includes natural changes in the brain, inducing sleep; the action of narcotic, hypnotic, or (some) hallucinogenic pharmaceutical agents; violent episodes such as epilepsy, narcolepsy or physical concussion. The fate of a participant, under these circumstances, is of considerable interest and the theory makes several predictions in this matter (most of which are confirmed). The possibilities of (b) are induced by sensory deprivation, confinement, isolation and the like and we predict that the evidence would be ambiguous (which it is; see, for example, Castaneda (1968, 1970), Huxley (1967) Lilly (1956, 1961)) unless the extent of the participant's actual L Processor are well defined; in which case, the results become much more predictable. This category of treatments is also ambiguously defined in itself, insofar as symbolic operations; for example, the overload experiments (Pask and Lewis 1962, Pask 1964) may be classed under the heading (b). The explicitly symbolic transformations of (c) are induced by hypnosis, relaxation and ritualistic exercises; some regarded as social phenomena. The effects induced by these symbolic acts are quite definite. The most cogent summary of this field is due to Clark (1969, 1972) who's own work on hypnosis and mystical states is both hard headed and beautiful.

6. A construction for a Self-Conscious Participant

Self-consciousness of A, with B, of R_i in L, means that A acts in the position of the external observer of Fig. 1, the modelling facility being the interface ι^0 between L processors α and β in

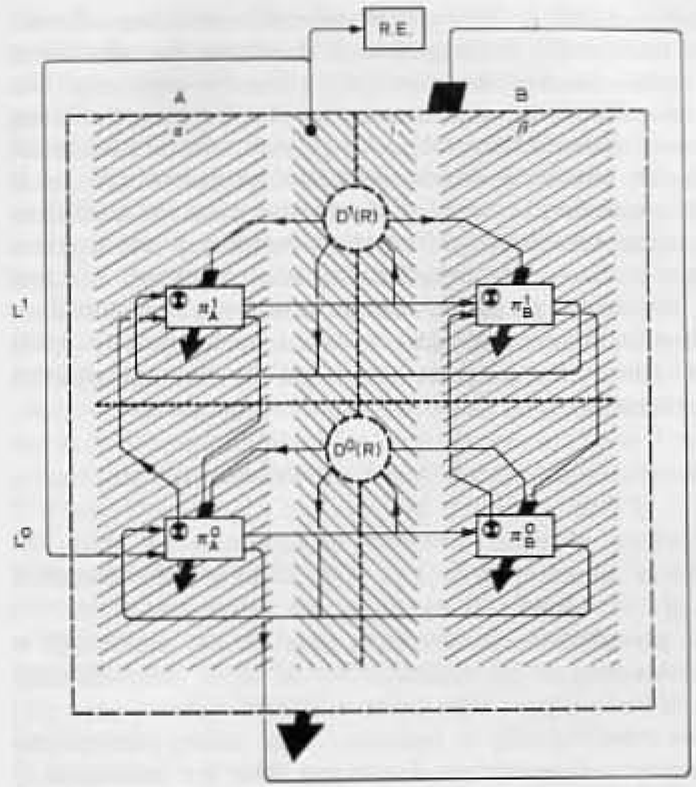

Icon 25.

which the model under execution is the L conversation between A, himself, and B. The construction is shown in Icon 25 and is subject to the restriction imposed upon Fig. 1 in order that an observer (here A) shall have the status of an external observer with respect to a conversation involving himself; namely, that A regards his modelling facility as one clocked. The restriction appears in Icon 25 as an identification between n and τ (written n → τ) and between t and T (written t → T). Because of this constraint, an occasion, n, advances the execution clock in A's modelling facility and also (by definition of an occasion) resets t (and thus T), if any modelling operation has been performed. In other words, self observation is restricted to the observation of alternative sets (any chain of explanation is truncated).

The form of self-consciousness permitted by Icon 25 is thus very narrow (it does however, encompass the requirement of

Howard (1966, 1973)). This form of self-conscious, though restricted, is not trivial. For example, if A adopts the calculating rules of a statistician, he can construct a macrotheory over the alternative sets or choice sets erected by his limited modelling operations and thus obtain estimates of his own doubt/belief regarding R_i, in terms of direct uncertainty/information (as in Section 1.3.3) or indirect estimates of selective work (as in Section 1.3.2). The actual transfer of information between A and B across the interface (alias, the synchronisation, because of this interaction, of otherwise asynchronous processes in α and β) is thus identified in the construction, with the participant A's (B's) doubt/belief about R_i as it is estimated by an uncertainty/information measure.

7. A Construction for Limited Veridical Statements

Add a further external observer (communicating with the existing external observer in L*) to Fig. 1. Many other external observers might be added (for example, the entire peer group of a science) but, provided the observers in question see themselves as occupying roles such as "scientist", two of them are sufficient; they are two P Individuals (with the role "scientist").

Engulf the modified Fig. 1 in Icon 25 by calling the original external observer a P Individual A and the other a P Individual B. The observers in question are no longer external; but they can act as though they were (that is, play the role of "scientist") and the price paid for doing so is to restrict their transactions in L* to transactions in L which must be constrained by the further caveats of the last section ($n \rightarrow \tau$ and $t \rightarrow T$). Under these circumstances, it is certainly possible for A to demarcate an understanding between A (himself) and B (the rest of the scientific establishment) about some R_i initially ostended at occasion n. In particular, A can agree with B about R_i which is the *L metaphor* of a strict conversation (used earlier in place of the usual *proposition*). As pointed out in Chapter 5, the L metaphor denotes a material analogy but, being a statement of agreement, does not in the ordinary sense, have a verdical truth value.

As promised in Chapter 5, the idea of verdical truth is retrieved, in a very restricted form, by the construction just delineated. If A and B agree about R_i in a strict L conversation then A makes no direct comment upon the L Processor in which the self-explaining

or self-replicating procedures responsible for his agreement are executed. In other words, participant A takes the L Processors α and β for granted, and does not mention them (which is no more nor less than an extreme form of the personal orientation noted in Section 5). But A does mention α and β if he observes himself. In doing so, he asserts it possible to carry out the construction just described because there actually exist L processors α and β that execute the procedures entailed by the L metaphor. If so, the material analogy which the metaphor denotes is true, in the sense that it holds good in α, β, ι^0. Thus any "verdical" statement involves a component to do with agreement or consensus about R_i, together with a statement that the participant (A) exists as an observer. However, A's statements are restricted. In order to play the role of external observer he must restrict his self referential L* statements (those involving the exterior loop in Icon 25) to causal L transactions with a one clocked modelling facility.

8. Broader Interpretations

The icons are drawn with the perspective of an external observer who looks at an objective and impersonalised entity it (his modelling facility) and who engages only in causal transactions with it. Because of that, any more general construction of self-observation, to represent a less blinkered "scientist" for example, abrades the icon scheme until it ultimately evaporates.

For example, if the $n \rightarrow \tau$, $t \rightarrow T$, restriction is relaxed, A has a many clocked modelling facility to accommodate the concurrent execution of a Fuzzy algorithm and Icon 25 could be immersed in Icon 4. The construction is an obvious way to represent the case in which A is "self-aware" and uses his brain and body as a modelling facility. There is nothing wrong with the fact of the matter; unfortunately, the icon becomes meaningless. In general, this occurs if L becomes L* "rather than" L* being confined to a limited form of L. Since L* transactions are no longer distinguished; since the modelling operations may be provocative rather than causal, A is no longer distinct from B. Finally, the distinguished category of "external observer" is eliminated; all observation is participant and any statement is to some degree self-referential. We have a reflective theory.

None of these circumstances are counterfactual, remarkable or out of the ordinary. They characterise the everyday condition of man. But this condition, though it exists, is not susceptible to explanation in terms of a peculiarly stilted mode of discourse (the prerogative of external observers) in which substantive comments all refer to causal transactions with an impersonal it.

9. Mutual Models of Student and Teacher

This kind of intractability besets an attempt to expand the metagame in Icon 21 (where explanations are truncated to yield choice sets or sets of alternatives) to the general metagame of Icon 22, where the explanations are not truncated. If neither

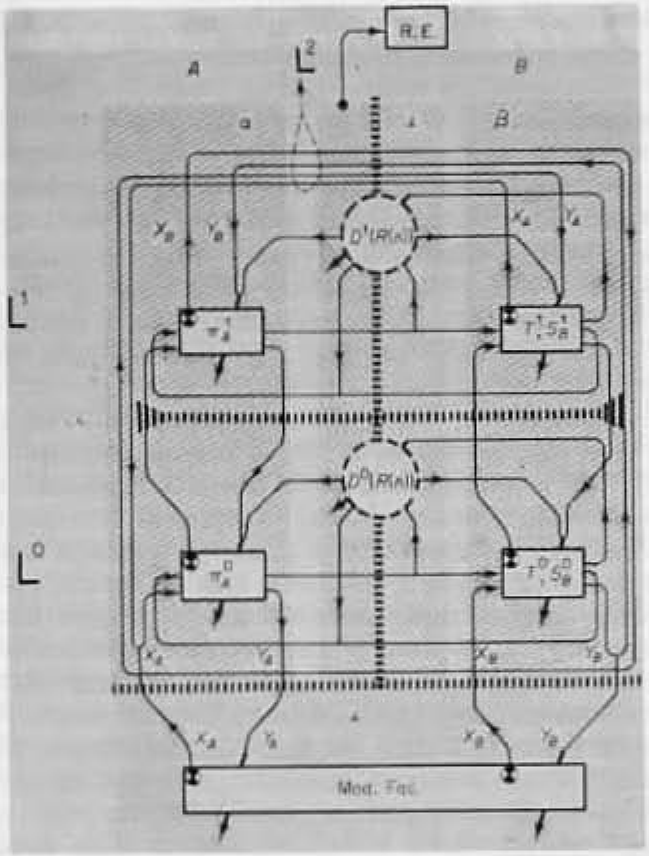

Icon 26.

participant pretends to be an external observer, the construction is possible. But the level of the chief model making loop (L^2) is ambiguous (L^2 may be "above L^1" or "below L^0" or both). Some of the L-dialogue may be unstratified and the L-conversation is no longer necessarily strict. The modelling operations are Fuzzy algorithms, used to construct A's models of B and B's models of A. If (as shown) A and B are executed in L Processors α and β the fuzzy modelling operations are carried out concurrently (not merely simulated).

A slightly specialised form of this paradigm has been designed for construction in the metal using the FRIM measures (Icon 24) as regulating variables. The special form is shown in Icon 26 and consists in a genuine tutor heuristic $B = \langle T^1, S^1_B; T^0, S^0_B \rangle$ where $\langle S^1_B, S^0_B \rangle$ is (as before) the support for A in the domain $\langle D^1 (R(n))\ D^0 (R(n)) \rangle$, and $\langle T^1, T^0 \rangle$ is a procedure for modelling A (in the broad sense, outlined above) in the context of the conversational domain. B's dominance as a teacher is secured by giving B overriding control of the evaluation of $\langle D^1(R(n)), D^0(R(n)) \rangle$, which is signified by the parametric operators penetrating their inscription at the common interface.

10. Embodiment in an Interface Processor

It is known that $\langle S^1_B, S^0_B \rangle$ can be realised in an interface processor that is not an L Processor (the cognitive reflector of the CET heuristic). The question of whether or not $\langle T^1, T^0 \rangle$ can be realised in this fashion is undecided and closely related to two pressing questions of "artificial intelligence"; namely, "can all of a P Individual be executed in an interface processor" and "can a procedure representing a dramatic character or personality (in the current nomenclature, a P Individual) be executed in an interface processor?"

The argument so far, which is taken up in the next volume, strongly suggests that neither of these two questions can be answered in the affirmative. But this does not mean that P Individuals can only be executed in brains. As hinted earlier in the discussion, inanimate L Processors can be constructed, and the only issue at stake is whether they compete with brains in terms of cost effectiveness.

Icon 26 is readily expanded by adding further participants $A = (\pi_A^1, \pi_A^0)$; say A_1 A_2 ..... each initially attached to a brain, say α_1, α_2 ..... The expanded icon represents a many person application of CASTE. Heuristic B is still dominant. The support $< S_B^1, S_B^0 >$ is executed at the interface (to which β is reduced). But, using a design based upon earlier work with small groups (Pask and Von Foerster 1961, Lewis and Pask 1966 and Pask and Lewis 1968) it is possible to distribute $< T^1, T^0 >$ for execution in some or all of the L Processors α_1, α_2 ..... This, at any rate, is a workable realisation and it is described in the next volume.

As a preview, not only is $< T^1, T^0 >$ distributed under execution but so are the P Individuals (A_1, A_2 ....). The student-to-student interactions that are catalysed in realising $< T^1, T^0 >$ ensure that A_i is executed to some extent in α_j and A_j in α_i. The main technical expedient is separate storage of the marker predicate assignments for each currently operating P Individual and replication of the support $< S_B^1, S_B^0 >$. Incidentally, the number of P Individuals is not necessarily the same as the number of students entering the system; coalitions are formed and fresh P Individuals may, in principle, be constructed as deviants or by co-operative combination.

In contrast to the idea of a "teaching machine", which was bandied around and mildly derrogated in Chapters 1 and 2, this arrangement is a "vehicle for driving through knowledge". That is, a student sees the entailment structure as a terrain through which he may progress, more or less freely, by learning and understanding. But he is subject to navigation rules imposed by the CET heuristic of (S_B^1 S_B^0) and boundary conditions that are imposed, through the macro measures discussed earlier in this chapter, perhaps augmented by FRIM indices. Further, the student expresses his desired direction and rate of motion in terms of macro variables. CASTE, regarded as a physical processor, provides an M Individual which the student (qua P Individual) inhabits and uses as a vehicle to drive through the terrain of the entailment structure. Each student may be seen (and may see himself) as occupying a separate vehicle. But students converse with one another, as well as with the support $< S_B^1, S_B^0 >$. In fact, student transactions carry the burden of teaching. Students are permitted to converse, if they occupy the same neighbourhood in the terrain and if they have learning strategies that are compatible.

General applications of SST systems are reviewed in Lewis and Pask (1964) and in Pask (1974); a summary account of the gross findings is given in Pask (1966). Detailed descriptions appear in Pask (1958, 1960, 1961, 1963, 1965, 1969). Most of the experimental work referred to the acquisition of coding skills in which a subject is presented with a complex stimulus (a configuration of several illuminated signal lamps, for example) and is able to achieve a correct result if he makes an appropriate complex response (pressing a configuration of several response buttons, for example) within a fixed interval of Δt. The response is appropriate if and only if the buttons pressed satisfy a coding rule that relates the display (of lamps) to the set of response buttons; Δt and the intertrial interval are chosen so as to render the task learnable but impossible for any novice. In our more sophisticated experiments, we regulated Δt to fit the individual subjects and took care, when analysing the data, that such individual adjustments did not pervert the data, at any rate in a direction that favoured the hypothesis being tested.

Most of the experiments were controlled by a special purpose computer on which any kind of SST or adaptive system could be simulated; the display arrangements included means for delivering either simple or differential 'knowledge of results' for optionally cueing in respect of each or all response components and for changing and alternating the current rule. The work is chiefly reported in Pask and Lewis (1967, 1968).

Salient conditions and results were as follows (each experimental subgroup has between 10 and 20 subjects).

(1) Using a single rule, and simplifying the problem (of realising the rule for a randomly selected sequence of lamp configurations within Δt) by varying the lamp number and by an auxiliary cueing adjustment, it was possible to confirm the result cited in the text. Learning rate is maximised by setting ξ so that problems are maximally difficult whilst remaining intelligible; for values of ξ spaced $1/10\ \rho$ apart, the predicated trend is significant at the 0.1% level. The auxiliary cueing is introduced to avoid a pathological inversion, noted by Van Der Veldt and cited in Woodworth